VIRTUALIZATION PERFORMANCE STUDY

ClearCube Engineering Systems Integration and Testing White Paper

Spring 2007

Table of Contents

8834 Capital of Texas Highway North Austin, Texas 78759 **Voice** 512.652.3500 **Fax** 512.652.3501 **Toll Free** 866-652-3500

sales@clearcube.com **www.clearcube.com**

INTRODUCTION

The term virtualization has been widely used since the 1960s or earlier, and has been applied to many different aspects and scopes of computing — from entire computer systems to individual capabilities or components. The common theme of all virtualization technologies is the hiding of technical detail.

According to VMware, the leader in virtual infrastructure software for industry-standard systems, virtualization is an abstraction layer that allows multiple virtual machines, with heterogeneous operating systems to run in isolation, side-by-side on the same physical machine.

ClearCube Technology offers its customers Sentral™ Software to manage virtualized desktop environments and delivers, through its blade computers, a flexible platform for virtualization.

As of yet, no industry standard for a performance measurement exists, so ClearCube recently conducted a study on the ways that typical users might interact with virtualized sessions on a PC blade.

This paper provides the performance results of that study and includes guidelines to install and configure virtualization products from VMware in similar environments in your corporate installations. In the study, ClearCube answers often-raised questions, including “how many virtual machines (VMs) can be installed in a standard configuration?” The information provided can be used as a baseline to determine how many VMs you can reasonably utilize in your installation environment. Later in the paper, you will find detailed performance thresholds for VMware ESX and GSX. ClearCube determined these thresholds during stress tests of numerous configurations.

1. Performance Summaries

The following charts provide the study results. Throughout the paper, we will refer to these results.

In every configuration in this report, each virtual machine (VM) was configured with the following PC blade configurations:

- Microsoft XP-SP2 Operating System
- Microsoft Office 2003
- 512 MB RAM (assigned to the VM)
- 4 GB of disk space

A1010 base configuration:

- Windows Server 2003 as Host OS with Application Server configured “on”
- VMware Virtual Server 1.0.1 (aka GSX) providing the virtualization layer
- Virtual machines; each VM OS image is stored on the A1010’s hard drive

GB RAM	Maximum Virtual Machines
2	3
3	5
4	6

Table 1. Performance Results for A1010 + VMware Virtual Server 1.0.1

8834 Capital of Texas Highway North Austin, Texas 78759 **Voice** 512.652.3500 **Fax** 512.652.3501 **Toll Free** 866-652-3500
sales@clearcube.com www.clearcube.com

Maximum Virtual Machines		
GB RAM	Good Performance	Maximum Load
4	7	8
6	12	14
8	15	17

Table 2. Performance Results for R2200 + ESX

R2200 base configuration

- VI-3 (ESX) as virtualization hypervisor
- Dual 3.2 GHz Intel Xeon CPUs
- NFS server on an R1300 with Suse Linux v10 – serving as a NAS
- VMs; each VM O.S. image is stored on the NAS

2. Interpreting the Performance Results

The performance summaries shown in Tables 1 and 2 communicate the recommended maximum number of VMs that can be used to obtain reasonable performance. In this performance study, we evaluated three primary performance areas:

- CPU utilization
- Memory utilization
- I/O utilization

Typically, these three components are the main stress points in computing environments. During this study, however, the I/O utilization did not seem to be a factor in any of the configurations. (We will discuss this further in the R2200 configuration, which used a unique configuration for Internet Protocol-based storage.)

Now, we'll look more closely at the behavior of each main configuration starting with the A1010.

a. A1010 + VMware Virtual Server (GSX) Configuration

Consider an A1010 blade with 3 GB of RAM. This configuration can support five VMs. Since VMware Server 1.0.1 does not have a sophisticated memory-handling routine, it is not possible to add a sixth VM to the 3 GB RAM configuration and still achieve reasonable performance. In this case the memory utilization is:

GB RAM/VM	Total GB	Purpose
0.5	0.5	Required by the host OS and the VMware Server kernel
0.5	2.5	Required for 5 x 512MB RAM used by each VM
	3.0	**Total Memory Required**

Table 3. Memory Requirement for A1010 Virtual Machines

8834 Capital of Texas Highway North Austin, Texas 78759 **Voice** 512.652.3500 **Fax** 512.652.3501 **Toll Free** 866-652-3500
sales@clearcube.com www.clearcube.com

In the A1010 example with 3 GB RAM, if a sixth VM is configured and started, the whole system slows to a crawl. There is no memory sharing between the VMs. The VMware kernel + MS Server 2003 requires the full 512 MB of RAM in order to properly accomplish its task. You cannot "over-subscribe" the memory. What you install and configure is what you get!

If a sixth VM is installed, the system will try to accommodate the new VM by temporarily swapping out the allocated memory for one of the first five VMs. This faulty configuration will temporarily turn off one VM (an office user); the memory is used for the recently added sixth VM. This is not a reasonable configuration.

b. R2200 + VI-3 (ESX) Configuration

For this study, the R2200 was configured as shown in the following figure. ESX is typically setup to run on one machine and the storage for the VMs is setup to run on IP-based storage. VMs are accessed over a standard Ethernet switch. Here, NFS was setup using Suse Linux on an R1300. The local hard disk on the R1300 acted as a NAS (Network Attached Storage) device. There was no noticeable I/O degradation during any of the testing with ESX, and a subsequent set of tests using a commercial iSCSI SAN product did not visibly improve I/O performance.

Consider an R2200 blade with 6 GB of RAM. This configuration can easily support 12 VMs in a typical office worker situation. Using a sophisticated memory algorithm, ESX allows VMs to "over-subscribe" the physical memory in a system. ESX will use approximately 512 MB of RAM to run the ESX kernel. If each of the remaining VMs used 0.5 GB RAM as they did with VMware Server 1.0.1, the overall utilization would be 6.5 GB of RAM.

In this R2200 example, 12 VMs are able to run because of the over-subscribed memory. ESX evaluates the memory usage pattern of each VM. If the ESX kernel finds two exact copies of the same content (i.e. Microsoft Word application) in memory, it swaps out one of those copies. This frees memory to be used by other applications or VMs.

As a result, if each of the 12 VMs is running the exact same copy of Microsoft Word, only one copy of the application stays resident in memory and the remaining 11 copies are swapped out of memory. (The content of the 12 Word documents will remain in memory since each document is likely to be unique.)

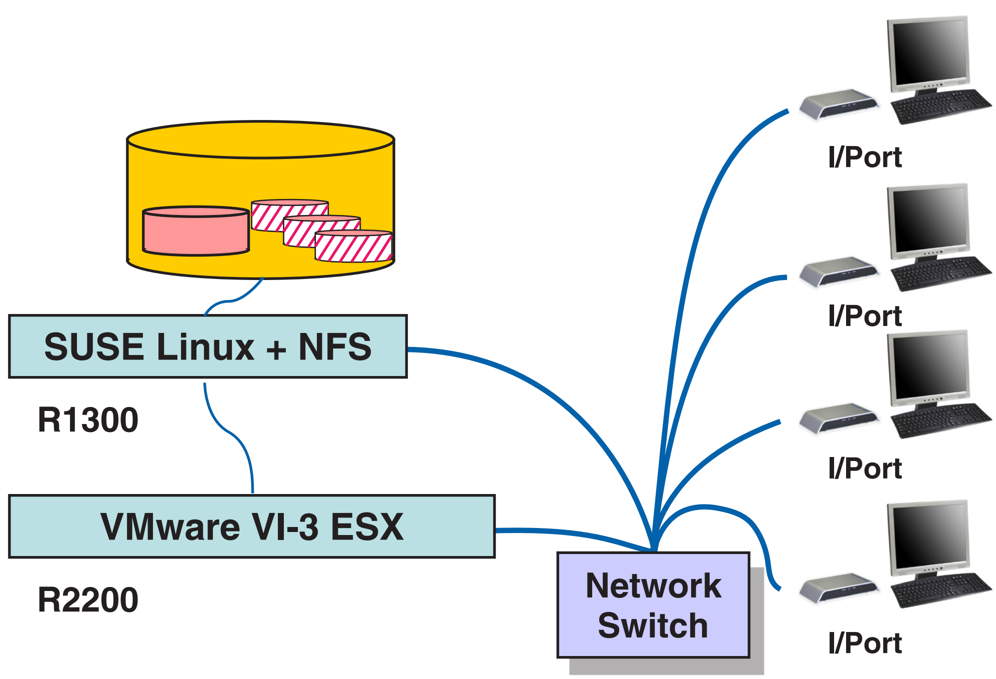

Fig 1. R2200 + VI-3(ESX) Configuration

8834 Capital of Texas Highway North Austin, Texas 78759 **Voice** 512.652.3500 **Fax** 512.652.3501 **Toll Free** 866-652-3500
sales@clearcube.com www.clearcube.com

i. ESX optimization

This ESX memory pattern matching and swapping algorithm improves over time on an individual ESX installation. As users access and use the same applications, ESX will more frequently swap out the common applications from memory. As the users continue to access the system with similar applications over time, overall system memory utilization will decrease, resulting in improved system performance over time. Within 24 hours of system usage, the ESX kernel has evaluated and balanced overall memory usage to derive an optimized configuration. Although numerous CPU cycles are required to achieve this efficiency, the feature works consistently to enhance system performance over time.

ii. ESX memory utilization considerations

Unfortunately this overall system efficiency is “reset” once the main ESX server is rebooted. Because the memory-matching algorithm is dynamic, ESX cannot predict and maintain memory utilization footprints between reboots. Once a server reboot occurs, the ESX kernel will reevaluate memory utilization and start swapping out memory as it recognizes common memory patterns making it possible to “over-subscribe” the memory (i.e. 12 VMs in a 6.0 GB RAM system).

This over-subscription takes time to become effective. If a 13th VM is added, the overall system slows to accommodate the new VM. Similarly, a 14th VM can be added, but the overall system is slow for a period of time as the ESX kernel makes memory adjustments. For an R2200 with 6 GB RAM, a configuration with 12 VMs will load and run efficiently after a reasonable period of time following a reboot. The amount of time it takes to readjust varies based on system load in the VMs. If all the VMs are quiescent, the memory adjustments to add an additional VM can occur in less than 20 minutes. If all the VMs are fairly active, then full memory readjustment may require up to 36 hours. As more VMs are configured into a system, the reboot readjustment period will generally take proportionately longer time periods.

This point should be heavily weighed when evaluating new installations of ESX in a corporate environment. If you have very stable power, rarely experience power outages, and have similar workloads, you may consider establishing a configuration with the maximum VMs for a given memory configuration (i.e. 14 VMs for 6 GB). However, if you reboot your ESX server during work hours, users on that ESX server may experience up to 30 minutes of outage as the system reboots and re-establishes the basic memory efficiencies. Configuring maximum VMs per server should be considered only if you have a very stable environment and perform hardware maintenance on weekends.

GB RAM/VM	Total GB	Purpose
0.5	0.5	Used by VMware ESX3
0.5	6.0	Required for 12 x 512MB RAM used by each VM
	6.5	Theoretical memory usage by the system
	6.0	**Actual memory used for the whole system**

Table 4. Memory Requirement for R2200 Virtual Machines

8834 Capital of Texas Highway North Austin, Texas 78759 **Voice** 512.652.3500 **Fax** 512.652.3501 **Toll Free** 866-652-3500
sales@clearcube.com www.clearcube.com

c. CPU Bound Versus Memory Bound

Reviewing the results in Tables 1 and 2, it is helpful to summarize which system component is limiting the configuration.

Memory-bound configurations have memory as the limiting component. If the system had more memory, more VMs could be added. CPU-bound configurations have the CPU(s) as the limiting component. If it were possible to use a faster CPU or add more CPUs, more VMs could be added. The following table shows the tested configurations.

Model	GB RAM	Recommended Maximum VMs	Memory Bound	CPU Bound
A1010	2GB	3	Yes	No
	3GB	5	Yes	No
	4GB	6	Yes	No
R2200	4GB	7	Yes	No
	6GB	12	Yes	No
	8GB	15	No	Yes

Table 5. Memory-Bound vs. CPU-Bound Considerations

In the table, the majority of the configurations are memory bound. For the A1010, as more memory is added to the 2 GB configuration, the system can support more VMs. When the system is fully configured with 4 GB, memory is still the limitation. In the A1010, CPU is not a limiting factor; it can easily handle 6 VMs with a typical office-worker load.

In the R2200, the 4GB configuration is memory bound. The two 3.2 GHz Xeon CPUs easily handle the load of seven VMs. As 2 GB is added to the configuration, memory is still the limiting factor and the CPUs are not used to capacity. However, as the R2200 reaches its maximum limit at 8 GB with 15 VMs, memory is still a factor and the two CPUs are running at a high capacity. At this point the CPUs are the limiting factor. (Find the actual memory and CPU limitation measurements at the end of this paper.) Table 5 shows the maximized CPU threshold in the R2200 with 8 GB and 15 or more VMs. Beyond 15 VMs, the CPUs become a limiting factor. If you add more VMs, system responsiveness is definitely decreased.

8834 Capital of Texas Highway North Austin, Texas 78759 **Voice** 512.652.3500 **Fax** 512.652.3501 **Toll Free** 866-652-3500
sales@clearcube.com www.clearcube.com

3. Performance Methodology Details

To make this paper applicable to many types of businesses, ClearCube simulated and measured performance in different configurations. This section communicates those results.

a. Simulating an Office Worker in a Corporate Environment

A standard baseline for most corporate configurations is an office worker who interacts with different types of documents throughout the workday. While some virtualization studies have focused on the use of Microsoft Word as the sole application for office use, this may be too simple an approach. In this study, we simulated a standard workload that uses multiple Microsoft Office applications. This load more realistically simulates a typical office worker's interaction and load on a system. Performance loads used include:

- Microsoft Word
- Microsoft PowerPoint
- Microsoft Excel

To make the performance measurements effective and repeatable, a scripting tool simulated a user interacting with the computer. A script, written to automatically create and edit documents, accomplished the following:

- Create a document
- Add content to the document at the speed of a normal person typing and using a mouse to interact with the system
- Edit the document to make changes
- Save and close the document
- Repeat

The script created multiple documents of each type (Word, PowerPoint, and Excel). Each document was created and left open while other documents were subsequently created and edited. This simulated a user interacting with different documents, and leaving some of them open for use later in their work session. After the multiple document editing was complete, the script closed all documents.

The script was setup to run multiple times until cancelled. In this way, the script was running on multiple VMs and on each VM simultaneously.

b. Performance Measurement and System Configuration

To simulate a full set of VMs on a fully loaded system, it is insufficient to simply install the software and run the performance test only on the VMware console. It is necessary to install a fully configured system as shown in Fig 2. ClearCube used I/Ports for the user component to ensure the IP network traffic emulated people typing and interacting with the entire system.

Consider the configuration of the A1010 blade with 3 GB RAM. An I/Port was configured and established connection with each VM during the configuration testing. To determine how many VMs were possible on this system, the performance was measured successively with the following VMs and I/Ports added:

- One VM using one I/Port; no other I/Ports connected to the system
- Two VMs using two I/Ports (one VM per I/Port); no other I/Ports connected
- Three VMs using three I/Ports (one VM per I/Port); no other I/Ports connected
- Four VMs using four I/Ports (one VM per I/Port); no other I/Ports connected
- Five VMs using five I/Ports (one VM per I/Port); no other I/Ports connected
- Six VMs using six I/Ports (one VM per I/Port)

This testing methodology is represented in the following figure.

Performance measurements were taken for each

8834 Capital of Texas Highway North Austin, Texas 78759 **Voice** 512.652.3500 **Fax** 512.652.3501 **Toll Free** 866-652-3500

sales@clearcube.com www.clearcube.com

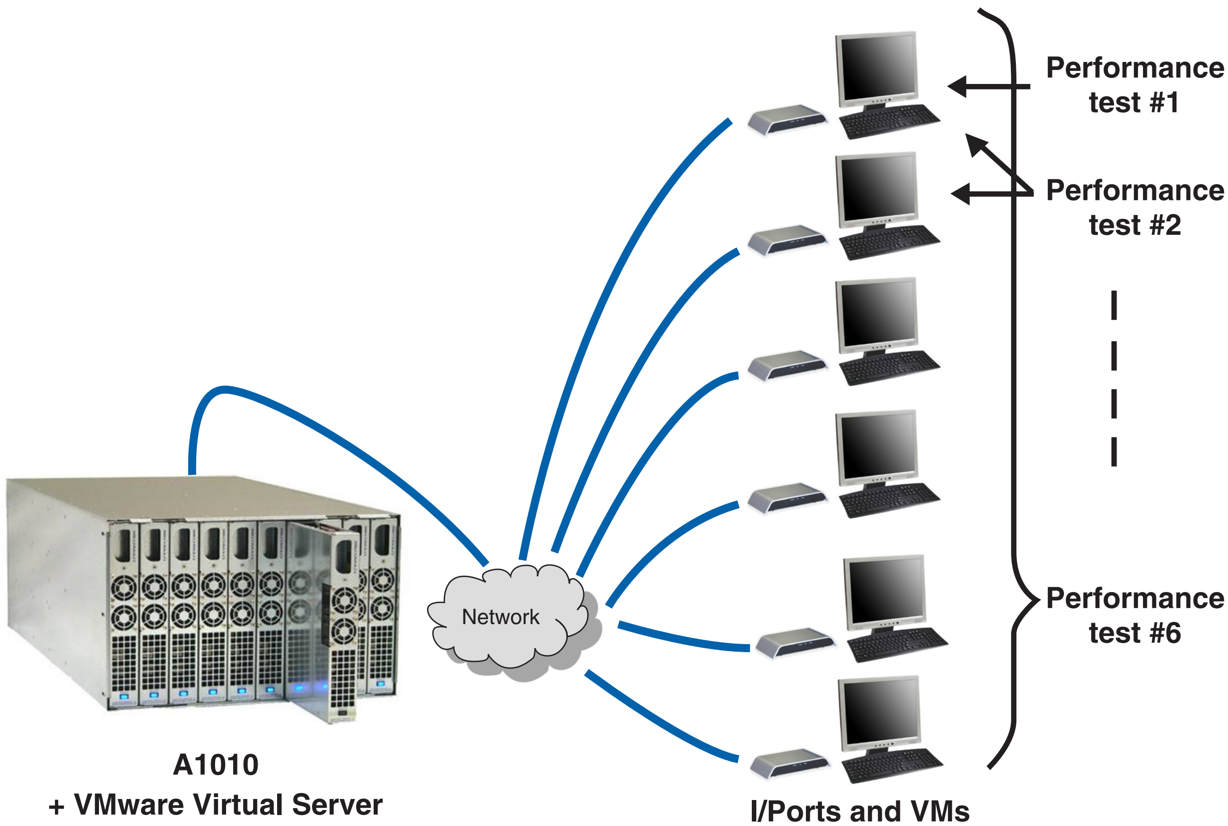

Fig 2. A1010 + VMware Virtual Server(GSX)

individual configuration. For example on the A1010 configured with 3 GB of memory, with five VMs, each user's simulated performance was fine. But when the sixth VM was added to this memory configuration, the system was unable to support the load.

4. Standards for Virtualization Performance Testing

Currently, there is no industry standard for performance measurement. While at least three proposed standards exist, none have been ratified as a formal industry-wide specification.

The Standard Performance Evaluation Corporation (SPEC) has formed a working group to study the virtualization performance testing and to develop a measurement technique. The SPEC working group is young, created in October 2006, and has not produced the standard or the methodology. View SPEC's most recent information at: http://www.spec.org/specvirtualization/

VMware has conducted some research and has proposed the VMware Benchmark. The benchmark, created by VMware employees, is not VMware specific and is targeted to become an industry benchmark. VMware is working with SPEC and SPEC-member companies to standardize this effort. As of February 2007, the VMware benchmark was in beta. See www.vmware.com for more details.

IBM and Intel have also proposed a performance benchmark christened as Vconsolidate for SPEC. In the past, performance benchmarking from these two companies has typically targeted very high-end "ultimate" server configurations.

See:

http://www.intel.com/technology/itj/2006/v10i3/7-benchmarking/6-vconsolidate.htm or

http://www.technologynewsdaily.com/node/5477

for more details.

8834 Capital of Texas Highway North Austin, Texas 78759 **Voice** 512.652.3500 **Fax** 512.652.3501 **Toll Free** 866-652-3500
sales@clearcube.com www.clearcube.com

5. Conclusions

This paper shows ClearCube performance results for the A1010 and R2200 platforms using VMware products. To reach its conclusions about performance, ClearCube evaluated the way a typical office user would interact with his virtualized session on a PC blade. The company then developed a scripted approach to simulating that user experience. From that methodology, ClearCube was able to determine a set of recommended configurations.

In your corporate PC blade deployments, you can use these results to determine a similar set of configurations. If you have users who are more task-oriented, they may stress the system less, allowing you to deploy more VMs on each configuration. Alternatively, you may have a set of users who place more demand on the virtualized environments. In this case, we recommend assigning fewer VMs to each configuration in order to deliver a better user experience.

In any configuration, ClearCube PC blades provide a very capable and flexible solution platform for deploying virtualization in your corporate environment.

8834 Capital of Texas Highway North Austin, Texas 78759 **Voice** 512.652.3500 **Fax** 512.652.3501 **Toll Free** 866-652-3500
sales@clearcube.com www.clearcube.com

Appendix A - Measured Test Results

The table below summarizes the results of the virtual desktop benchmarking that was conducted using the methodology described in section 3 above.

For a dual processor R2200 blade, 6 GB of RAM provides the best price/performance ratio for virtualized desktops. Because of the sophisticated VMware ESX3 memory management, you can lightly "over-subscribe" physical memory and still have this configuration easily support 12 VMs in a typical office worker situation.

	4 GB of RAM		6 GB of RAM		8 GB of RAM	
No of VMs	CPU usage %	Mem Used in GB	CPU usage %	Mem Used in GB	CPU usage %	Mem Used in GB
1 VM	7	1.0	8	1.0	8	1.0
2 VM	16	1.4	15	1.4	16	1.5
3 VM	21	1.8	22	1.8	22	1.7
4 VM	30	2.5	28	2.1	30	2.0
5 VM	35	3.0	35	2.3	35	2.3
6 VM	43	3.4	40	2.6	43	2.6
7 VM	49	3.6	49	2.8	49	3.0
8 VM	56	3.7	54	3.3	55	3.3
9 VM	59	3.7	59	3.7	60	3.5
10 VM			65	3.8	65	3.9
11 VM			69	4.3	69	4.1
12 VM			80	4.6	70	4.2
13 VM			85	4.9	72	4.7
14 VM					84	5.0
15 VM					88	5.2
16 VM					87	5.5
17 VM					90	6.0

Table 6. R2200 PC Blade - VMs Supported Versus Memory

8834 Capital of Texas Highway North Austin, Texas 78759 **Voice** 512.652.3500 **Fax** 512.652.3501 **Toll Free** 866-652-3500

sales@clearcube.com www.clearcube.com

AUSTIN, TEXAS WASHINGTON, D.C. NEW YORK LONDON MUNICH TOKYO